GAMES WITH CONTINUOUS, CONVX PAY-OFF

H. F. Bohnenblust, S. Karlin, L. S. Shapley

P-66

Revised

12 August 1949

20 p $1.00 he
$0.50 mf

GAMES WITH CONTINUOUS, CONVEX PAY-OFF

H. F. Bohnenblust, S. Karlin, L. S. Shapley

I. Background.

In the "normal form" of a two-person, zero-sum game, as the theory has been set forth by von Neumann [3], there are just two moves. They are the choices of strategy, made simultaneously by each player. One player is then required to pay to the other an amount (positive or negative) determined by the pay-off function, which is a function only of the strategy-choices. The theory is best known at present for games in which the number of strategies available to each player is finite. This article will explore a rather special class of games in which the strategies of one player form a compact and convex region B of finite-dimensional Euclidean space, while those of the other form an arbitrary set A.

In general, equality may or may not hold in

$$(1) \qquad \sup_{x \in A} \inf_{y \in B} M(x, y) \leq \inf_{y \in B} \sup_{x \in A} M(x, y).$$

Intuitively, there may be a gap between what the x-player's best "safe" strategy guarantees to him and what his opponent's best "safe" strategy prevents him from obtaining. When equality does not hold, the typical procedure of game theory is to replace the choosing of a strategy by the choosing of a probability distribution over the whole set of strategies. Thus, the player entrusts the

task of playing the game to a machine which makes random decisions, and contents himself with controlling its probable behavior to maximize his probable gain. Such a probability distribution is called a _mixed strategy_, and its _order_ is the number of points in the spectrum of the distribution. (That is, the order is infinite unless a finite set of strategies exists which is chosen with probability one; in that case the order is the number of strategies which are chosen with positive probability.) A _pure strategy_ is a mixed strategy of order one.

The game on the unit square will illustrate the use of mixed strategies without the inconvenient notation that general sets A and B would entail. Let A = B be the closed one-dimensional unit interval $[0, 1]$. Then, corresponding to (1) is the inequality

$$(2) \qquad \sup_{F \in \mathcal{D}} \inf_{y \in B} \int_{-0}^{1} M(x, y)dF(x) \leq \inf_{G \in \mathcal{D}} \sup_{x \in A} \int_{-0}^{1} M(x, y)dG(y),$$

where $\mathcal{D}$ is the set of all (cumulative) probability distributions on $[0, 1]$. ($F \in \mathcal{D}$ if and only if (i) $x < x_1$ implies $F(x) \leq F(x_1)$, (ii) $x < 0$ implies $F(x) = 0$, (iii) $x \geq 1$ implies $F(x) = 1$, (iv) F is continuous to the right.) Under quite general conditions, equality holds in the expression exemplified by (2), while not necessarily holding in (1). When it does hold, the number thereby defined is termed the _value of the game_. A distribution which achieves that value is termed an _optimal mixed strategy_ (_o.m.st._) for the player in question. Any pair of o.m.st. is termed a _solution_ of the game. A game may in some cases have a value without having a solution.

II. Summary and discussion of results.

A function φ is said to be convex if and only if, for any λ_1 and λ_2 satisfying

$$(3) \qquad 0 \leq \lambda_1 = 1 - \lambda_2 \leq 1,$$

the inequality

$$(4) \qquad \lambda_1 \varphi(x_1) + \lambda_2 \varphi(x_2) \geq \varphi(\lambda_1 x_1 + \lambda_2 x_2)$$

holds whenever all three terms are defined. It is strictly convex if, in addition, $x_1 \neq x_2$ and $\lambda_1 \lambda_2 \neq 0$ always imply the strict inequality in (4). The present paper deals with games in which the pay—off $M(x, y)$ is, for every x in A, a continuous convex function of y. Continuity in y and compactness of B are enough to assure the existence of a value, as has been shown by Wald [6]. Convexity in y further assures the existence of an optimal pure strategy for the y—player, that is, an o.m.st. of order one. The central result of the present paper is that the x—player must have an o.m.st. of order at most n + 1, where n is the dimension of B. Moreover, if the y—player has a p—dimensional set of o.m.st. of order one, then the x—player has an o.m.st. of order at most n − p + 1.

Without convexity the solutions, even of games on the unit square, may be much more complicated. If M is a polynomial, Dresher has shown that o.m.st. of finite order exist for both players [2]. But Blackwell and Girshick have found a unit square

game with continuous pay—off in which the only o.m.st. for each player makes use of every strategy [2].

It might be worth—while to illustrate the way in which the results for convex games can be applied to other games. A <u>linear</u> function is of course convex, and the expected pay—off of a game is always linear in the mixed strategies. It follows that, in any game, if B is a finite set with m elements, then the x—player has an o.m.st. of order m or less. A more general statement is that if B can be subdivided into m closed, convex, non—overlapping components B_i, of dimension n_i, such that the pay—off is convex over each component, then the y—player has an o.m.st. of order at most m and the x—player one of order at most $m + \sum n_i$. The verification of either statement is accomplished by constructing an equivalent, convex game with an enlarged set B' of strategies for the y—player.

Symmetrically corresponding assertions obviously hold, here and throughout the paper, with concavity in x replacing convexity in y.

Thus consideration of convexity (concavity) is a handy tool for uncovering the existence of simple solutions in potentially complicated games. The question of computing such solutions when they exist will be discussed in $\S V$ of the present paper.

III. <u>Theorem on convex functions</u>.

Let B be a compact, convex region in an $(n - 1)$—dimensional space whose elements are denoted by y. A function f is <u>linear</u>

(non-homogenous) if $f\left(\sum \lambda_i y_i\right) = \sum \lambda_i f(y_i)$ when $\sum \lambda_i = 1$. The function $f(y) \equiv 1$, denoted by $\underset{\sim}{1}$, is linear. The linear functions form an n-dimensional linear space E. F will denote an element of the conjugate space E^*.

LEMMA 1.1. If $F(\underset{\sim}{1}) = 1$ there exists y such that $F(f) = f(y)$ for all f in E.

Proof: It suffices to show that the n equations $F(f_i) = f_i(y)$ have a solution in y for n linearly independent elements f_i of E. But one may take $f_1 = \underset{\sim}{1}$ and get an identity for the first equation. The remaining $n - 1$ equations, still independent, have a solution.

LEMMA 1.2. The set of all f which are non-negative over B forms a closed convex cone $P \subset E$, with vertex at the origin, containing $\underset{\sim}{1}$ in its interior. Moreover, the region over which $P(y) > 0$ is precisely B.

(The notation "P(y)" will mean "f(y) for all f in P"; "f(B)" will mean "f(y) for all y in B.")

Proof: The first part is obvious. For the latter, B and any y not in B can be separated; i.e., some f in E will have $f(y) < c \leq f(B)$. Then $f - c$ is in P and is negative for y.

LEMMA 1.3. Let Q be a compact convex set of E which does not intersect P. There exists y in B and δ such that $Q(y) \leq -\delta < 0$.

Proof: P and Q are separated by some F in E^*; that is, for some $\delta > 0$, $F(Q) + \delta \leq F(P)$. Since P is a cone with vertex at the origin and F(P) is bounded from below, F(P) must be non-negative. Since $\underset{\sim}{1}$ is in the interior of P, making $F(\underset{\sim}{1}) > 0$, F may be chosen

so that $F(\underset{\sim}{1}) = 1$. By Lemma 1.1 a y exists satisfying

$$q(y) + \delta \leq 0 \leq P(y),$$

while, by Lemma 1.2, y must be in the set B.

LEMMA 1.4. If $p_1, \cdots, p_m$ are points in an $(n-1)$-dimensional space, then any point in their convex is in a convex spanned by at most n of them.

Proof: Take a simplex S_m in $(m-1)$-dimensional space and a linear transformation mapping it on the given convex C, the vertices of S_m going into the points $\{p_i\}$. The inverse transformation maps each point p of C onto a plane $L(p)$, of dimension at least $m - n$, which intersects S_m. When a plane meets a simplex but not its boundary, the intersection is a point. Hence there is a simplicial face of S_m which intersects $L(p)$ in a point. Its dimension must be less than n, and its vertices obviously correspond to a subset of $\{p_i\}$ which spans p.

LEMMA 1.5. If $\sup_\alpha f_\alpha(B)$ is positive for a family of $\{f_\alpha\}$, then for suitable $\lambda_i \geq 0$, $\sum \lambda_i = 1$, and α_i, $i = 1, \cdots, n$, the function $f = \sum \lambda_i f_{\alpha_i}$ is in P; that is, $f(B) > 0$.

Proof: The Heine-Borel covering theorem permits one to work with a finite sub-family of $\{f_\alpha\}$, since $f_\alpha(y) > 0$ defines an open set. (This is the only use made of strict positiveness. The hypothesis might alternatively read "If $\sup_\alpha f_\alpha(B) \geq 0$ for a finite family ...". In this form Lemma 1.5 is equivalent to Ville's lemma [5].) The convex Q spanned by the finite sub-family must intersect

P, by Lemma 1.3. Since Q is bounded and P is not, some boundary point of Q lies in P. This point is on a polyhedral face of dimension at most $n - 1$. Lemma 1.4 now gives us the desired representation.

THEOREM 1. Let $\{\varphi_\alpha\}$ be a family of continuous convex functions defined over a compact, convex, $(n - 1)$-dimensional region B. Then $\sup_\alpha \varphi_\alpha(y)$ attains its minimum value c at some point of B; and, given any $\delta > 0$,

$$\sum_{i=1}^n \lambda_i \varphi_{\alpha_i}(B) \geq c - \delta,$$

for any suitable choice of α_i and $\lambda_i \geq 0$, $\sum \lambda_i = 1$.

Proof: Let

$$a > \inf_{y \in B} \sup_\alpha \varphi_\alpha = c.$$

The set of y in B with $\varphi_\alpha(y) \leq a$ is non-void, closed and convex, and decreases as a decreases. The intersection of all these sets is non-void, and any point in it satisfies the first part of the Theorem. For the second part, let $\{f_\beta\}$ be the family of linear functions with $\left[\varphi_\alpha - f_\beta\right](B) \geq 0$ for some α. This family contains all planes of support to all φ_α; therefore $\sup_\beta f_\beta = \sup_\alpha \varphi_\alpha$. Apply Lemma 1.5 to the family

$$\mathcal{f}_\delta = \left\{f_\beta - (c - \delta)\underset{m}{1}\right\}, \qquad (\delta > 0).$$

Each β_i so obtained corresponds to an α_i with $f_{\beta_i} \geq \varphi_{\alpha_i}$. These α_i and the λ_i of the Lemma provide the representation of the Theorem.

COROLLARY 1.1. If the (convex) set Y of points for which $\sup_\alpha \varphi_\alpha(y) = c$ has dimension p, then the number of functions φ_{α_i} required is at most n − p.

Proof: Take an (n − 1 − p)—dimensional cross section $B' \subset B$ perpendicular to Y and intersecting Y in an interior point y_0. Let d_1 be the distance from y_0 to the nearest boundary point of Y, and let d_2 be the diameter of B. Apply the Theorem to B' and $\delta' = \delta d_1/d_2$. The λ_i and α_i so obtained, i = 1, $\cdots$, n − p, must work for the original B and δ.

The following will be obtained in a somewhat different form in § VII and is put here for the sake of completeness.

COROLLARY 1.2. If Y is in the boundary of B, then the number of functions required is at most n − p − 1.

IV. Application to games.

Consider a bounded pay—off function M(x, y) where the choice x [y] of the maximizing [minimizing] player is taken from the set A [B]. M is continuous and convex in y for each x, and B is a compact, convex region in (n − 1)—dimensional Euclidean space. Let $Y \subseteq B$ denote the set of points which minimize $\sup_x M(x, y)$ and let p denote the dimension of Y. (Y is non—void, closed and convex.) Let I_z denote the pure strategy by which the point z is chosen with probability one. A mixed strategy will be called ε—effective if the value of the game is not more than ε better than the expected return guaranteed by the mixed strategy to its user. Thus, an o.m.st. is 0—effective.

THEOREM 2. The value of the game described is

$$c = \min_{y \in B} \sup_{x \in A} M(x, y).$$

For any $\varepsilon > 0$, there is an ε-effective mixed strategy for the x-player of the form

(5) $$F_0 = \sum_{i=1}^{n-p} \lambda_i I_{x_i} \qquad \left(\sum_{i=1}^{n-p} \lambda_i = 1, \lambda_i \geq 0 \right);$$

while all pure strategies I_{y_0} on some y_0 in Y are optimal for the y-player.

Proof: The Theorem is a direct consequence of Theorem 1 and Corollary 1.1.

COROLLARY 2.1. If in addition A is compact and M is continuous in x for each y, then some mixed strategy of the form (5) is optimal.

Proof: The added conditions make $\left\{ M(\alpha, y) \right\}$ a closed family, hence it is permissible to take $\varepsilon = 0$ in the Theorem.

COROLLARY 2.2. If, moreover, M is strictly convex in y for each x, then the y-player's o.m.st. is unique.

Proof: Using some fixed o.m.st. $\sum \lambda_i I_{x_i}$, define Y_ν as the set of y with

$$\sum \lambda_i M(x_i, y) < c + 1/\nu \qquad \nu = 1, 2, \cdots.$$

Let P denote the set function associated with any optimal y-strategy.

Then it is easily seen that $P(B - Y_\nu) = 0$ for any ν. But strict convexity implies that $\cap Y_\nu$ is a single point y_0; hence the o.m.st. I_{y_0} is unique.

V. Computation of the solution.

Suppose, to avoid the complication of $\mathcal{E}$-effective mixed strategies, that the conditions of Corollary 2.1 are met, so that the game has an o.m.st. of the form (5). The determination of the value of the game

$$c = \min_{y \in B} \; \max_{x \in A} \; M(x, y)$$

and of the sets

$$Y = \text{those } y \text{ for which } \max_{x \in A} M(x, y) = c,$$

$$X = \text{those } x \text{ for which, for all } y \text{ in } Y, \; M(x, y) = c,$$

must be considered a routine computation in the present discussion, since any difficulty here will have arisen from the nature of the unspecified set A.

To complete the solution it is sufficient (a) to discover a finite subset $X' = \left\{ x_i \right\}$ of X with

$$\min_{y \in B} \; \max_{x \in X'} \; M(x, y) = c$$

and then (b) to find weights λ_i for the x_i not more than n-p of which are actually positive, and which make $\sum \lambda_i I_{x_i}$ an o.m.st.

The process is best described geometrically.

Let y_0 designate a fixed interior point of Y, and let B' be some $(n - p - 1)$-dimensional cross section of B, meeting Y in precisely y_0. Each x of X describes a convex hyper-surface over B' which has one or more supporting hyper-planes at (y_0, c). Let S be a small sphere in B' with y_0 as center, and with each x associate the set S_x of points in S corresponding to the directions of steepest ascent of all the planes of support to $M(x, y)$ at (y_0, c). S_X will denote the union of the S_x for $x \in X$.[*]

The progress of the reduction may be traced through the following four statements which, for any fixed finite $X' \subseteq X$, are either all true or all false:

$$\min_{y \in B} \max_{x \in X'} M(x, y) = c,$$

$$\min_{y \in B'} \max_{x \in X'} M(x, y) = c,$$

(6)

$$\lim_{\rho \to 0} \min_{y \in S} \max_{x \in X'} M(x, y) = c, \text{ where } \rho \text{ is the radius of S,}$$

$$\min_{y \in S} \max_{y' \in S_{X'}} y \cdot y' \geq 0.$$

The inner product $y \cdot y'$ is taken relative to S as the unit sphere. Thus $y \cdot y'$ is the cosine of the angle $\measuredangle y y_0 y'$.

---

[*]If any plane is horizontal, the game is solved instantly, since the plane must correspond to a pure optimal x-strategy.

Two assumptions must be interjected here:

(7)   Y is not in the boundary of B,

(8) $\min\limits_{y \in S} \max\limits_{y' \in S_X} y \cdot y' = d > 0.$

Without the first, S would contain non-strategies. Without the other it becomes more difficult to show that the computation is finite. As will be seen in §VII, failure of either (7) or (8) actually reduces the order of the optimal x-strategy, thus simplifying the computation. Geometrically, (8) states that y_0 is interior to the convex in B' spanned by S_X.

To continue: select y_1' at pleasure from S_X and proceed by the recursive instructions $(k = 2, 3, \cdots)$:

(I_k) If

$$m_k = \min\limits_{y \in S} \max\limits_{i < k} y \cdot y_i' < 0,$$

then let y_k denote the (unique) y in S at which the minimum occurs; if $m_k \geq 0$, terminate the process.

(II_k) Let y_k' denote a point of S_X for which[*]

$$y_k \cdot y_k' = \max\limits_{y' \in S_X} y_k \cdot y'.$$

---

[*] Note that S_X is a closed set.

The iteration terminates after a finite number of steps. For if not, there would be $k_1 < k_2 < \cdots < k_\ell < \cdots$ for which the subsequences y_{k_ℓ} and y'_{k_ℓ} both converge. But then

$$\lim_{\ell \to \infty} m_{k_\ell} \geq \lim_{\ell \to \infty} y_{k_\ell} \cdot y'_{k_{\ell-1}} = \lim_{\ell \to \infty} y_{k_\ell} \cdot y'_{k_\ell} \geq d > 0$$

implies a finite termination after all.

The $\{y_1, \cdots, y_m\}$ so obtained leads back to a set $X' = \{x_1, \cdots, x_m\}$ for which the statements (6) are all true. Moreover a particular supporting plane $P_i(y)$ is denominated for each x_i. The supporting planes are distinct, but the x_i may not be.

The weights which solve the original game $M(x, y)$ will also solve the semi-discrete, linear game

$$P_i(y) = x_i \cdot y + c \qquad\qquad (i = 1, \cdots, m; y \in S),$$

and conversely. The y-player here does not have a pure o.m.st. since the point y_0 is denied him, but any convex combination of $y \in S$ giving y_0 will be optimal, by the linearity. The linearity moreover makes it sufficient to consider the equivalent, wholly discrete game

$$||P_{ij}|| = ||P_i(y_j)|| \qquad\qquad (i = 1, \cdots, m; j = 1, \cdots, n-p)$$

where the y_j are any $n-p$ points on the sphere S whose convex contains a neighborhood of the center, y_0. The y-player here has a unique o.m.st. with all weights positive. It follows (see [4]) that $m \geq n-p$

and that some n-p x n-p submatrix P' will have the property

$$\lambda_\nu = \frac{\sum\limits_{j} P_{ji_\nu}}{\sum\limits_{\nu}\sum\limits_{j} P_{ji_\nu}} \geq 0, \qquad \text{all } \nu = 1, 2, \cdots, n-p,$$

P_{ji_ν} being the cofactor in P' of $p_{i_\nu j}$. P' may be discovered by a finite inspection. Then

$$F_0(x) = \sum_{\nu=1}^{n-p} \lambda_\nu I_{x_{i_\nu}}(x)$$

is the desired o.m.st.

VI. The solution in one dimension.

A complete description of the solution in the case n = 2 will serve to point up the discussion of the preceding section. For definiteness, let M(x, y) be defined on the unit square A x B = [0, 1] x [0, 1], and let it be continuous in each variable with the cross section at each x a convex curve over B.[*]

Suppose first that $\max_x M(x, y)$ has a unique minimum c at a point y_0 interior to the interval B. Then the set X of convex curves passing through (y_0, c) will be the union of two sets X_ℓ and X_r, not necessarily disjoint, defined by:

---

[*] A. V. Martin collaborated with the authors in the original study of this case.

$$X_\ell = \text{those } x \in X \text{ with } M'_\ell(x, y_0) = \lim_{y \to y_0^-} \frac{M(x,y) - M(x,y_0)}{y - y_0} \leq 0,$$

$$X_r = \text{those } x \in X \text{ with } M'_r(x, y_0) = \lim_{y \to y_0^+} \frac{M(x,y) - M(x,y_0)}{y - y_0} \geq 0.$$

To obtain an optimal x-strategy, select any $x_1 \in X_\ell$, $x_2 \in X_r$ and assign non-negative weights λ_1 and $\lambda_2 = 1 - \lambda_1$ satisfying

$$\lambda_1 M'_\ell(x_1,y_0) + \lambda_2 M'_\ell(x_2,y_0) \leq 0 \leq \lambda_1 M'_r(x_1,y_0) + \lambda_2 M'_r(x_2,y_0) \, .$$

These weights will be precisely determined only[*] when $M(x_1,y)$ and $M(x_2,y)$ are actually differentiable at y_0. Otherwise there will be two extreme pairs of weights. Convex linear combinations of these extreme strategies, for all possible pairs x_1, $x_2 \in X_\ell$, X_r, will provide all o.m.st. of finite order for the x-player.

If $M'_\ell = 0$ for any $x \in X_\ell$, or $M'_r = 0$ for any $x \in X_r$, then that x represents a pure optimal strategy. (Cf. the second reduction discussed in § VII.)

The same formulation is valid for $y_0 = 0$ or $y_0 = 1$ if the convention $M'_\ell(x,0) = -\infty$, $M'_r(x,1 = +\infty$ is adopted. In these cases $X_r \subseteq X_\ell = X$ and $X_\ell \subseteq X_r = X$ respectively; hence o.m.st. of order one, among others, will be found (cf. Corollary 1.2 above and the first reduction of § VII).

The same formulation is also valid trivially if $\max_x M(x,y)$ has its minimum over an interval Y. If y_0 is any interior point of Y, then $M'_\ell(x,y_0) = M'_r(x,y_0) = 0$ for all $x \in X$. In this case all the extreme o.m.st. are pure. (This is the case $n = 2$, $p = 1$.)

EXAMPLE. Let $M(x,y) = f(y-x)$ in the unit square with $f''(u) > 0$ for $u \in [-1, 1]$. Suppose $f(-1) > f(0) < f(1)$, then the

_____
[*] The weights $(0,1)$ will be determined uniquely even without differentiability at y_0, provided that $M'_r(x_2, y_0) = 0 > M'_r(x_1,y_0)$; also the weights $(1,0)$ in the symmetric case.

equation $f(u) = f(u-1)$ has a unique solution $u = a$, $0 < a < 1$.
In the light of the preceding discussion the following results may
be stated:

(a) The value of the game is $f(a)$;

(b) The unique optimal y–strategy is I_a;

(c) The unique optimal x–strategy is $\alpha I_0 + (1-\alpha)I_1$, where α
is given by the equation $\alpha f'(a) + (1-\alpha)f'(a-1) = 0$.

If $f(-1) < f(0)$, or if $f(0) > f(1)$, then the unique optimal strategies
are I_0 for both players, or I_1 for both players, respectively, and
the value is $f(0)$. If $f(-1) = f(0)$ or $f(0) = f(1)$, or if one assumes
only that $f''(u) \geq 0$, the optimal strategies are in general not unique.

VII. Sharpening of the results.

The discussion of this section will dispose of assumptions (7)
and (8) of §V and the proof of Corollary 1.2 of §III, and concurrently
describe improved results for certain special situations.

First it may be remarked that, by using known properties of
discrete games (see [4], and [1] Theorem 1), two sharper conclusions
may be drawn from the matrix obtained in §V:

(i) every o.m.st. (of the discrete game) is a convex linear
combination of extreme o.m.st. of order n–p or less;

(ii) every strategy i participates in at least one such extreme
o.m.st. of the x–player.

Referred to the original game, (1) implies that all o.m.st. of finite
order may be put in terms of extreme o.m.st. of order n–p or less.

The construction of $\S V$, of course, does not lead to a complete set of finite o.m.st. (to say nothing of the infinite ones that can easily be shown to exist whenever X is infinite). But, in consequence of (ii) and the arbitrariness of y_1' in $\S V$, it will succeed in producing an extreme o.m.st. involving any one given x of X with positive weight.

Suppose now that Y is in the boundary of B, and hence that y_0 is in the boundary of B'. In order to contain the sphere S, B' must be enlarged. But if it is to become legal for the y-player to choose y from outside of B, it must also be made unprofitable, if the solution is not to be disrupted. Therefore, introduce a dummy strategy x_0 into the set A with pay-off

$$M(x_0, y) < c \qquad \text{interior to B,}$$

$$M(x_0, y) = c \qquad \text{on boundary of B,}$$

$$M(x_0, y) > c \qquad \text{exterior to B.}$$

This function may be made continuous and convex in y since B is a convex region. It may also be made arbitrarily "steep" as it crosses the boundary, making it unimportant whether or not it is actually possible to extend the other functions $M(x, y)$ convexly into the exterior of B. Now by the remark of the last paragraph an o.m.st. of order n-p or less may be found utilizing x_0 with weight $\lambda_0 > 0$. But the mixed strategy obtained by redistributing λ_0 among the other components, in proportion to their own weights, must be optimal in the

original game.[*] Therefore, <u>at</u> <u>least</u> <u>one</u> <u>of</u> <u>the</u> <u>extreme</u> <u>o.m.st.</u> <u>is</u> <u>of</u>
<u>order</u> n-p-1 <u>or</u> <u>less</u>. Removed from the games context this conclusion
becomes Corollary 1.2 of $\S$ III.

It might be remarked that a reduction of more than one — while
possible — can not be deduced in general from the hypothesis that Y
is situated in a lower-dimensional "corner" of the boundary of B.

To gather in the last loose end, suppose that assumption (8)
of $\S$ V does not hold. This would mean that along some directed line
in B' emanating from y_0 none of the set of supporting planes actually
increases. Equivalently, this would mean that the "bottom", Y, of
the hyper-surface $z = \sup_x M(x, y)$ is less extensive than the "bottom",
Y_L, of the envelope from above of the linear functions supporting
$M(x, y)$ at Y. The prescription for dealing with this situation,
should it occur, is simple: using Y_L in place of Y, define the
cross section B'_L and sphere S_L. Then, replacing (8) with

$$\min_{y \in S_L} \ \max_{y' \in S_{LX}} \ y \cdot y' = d > 0,$$

proceed with the computation. The results involving p = dim Y
will be replaced by stronger results involving p_L = dim Y_L < p. Thus,
unlike the boundary reduction detailed above, this case reduces <u>all</u>
the extreme o.m.st. to order $n-p_L$ or less.

Finally, it is clear that the two reductions just described act
independently, their effects being additive if both occur together.

---

[*] The formal proof is straightforward.

BIBLIOGRAPHY

1. H. F. Bohnenblust, S. Karlin, L. S. Shapley, Solutions of discrete, two-person games (to be published).

2. D. Blackwell, M. Dresher, M. A. Girshick, (private communication).

3. J. von Neumann, O. Morgenstern, Theory of Games and Economic Behavior, Ch. II, 2nd ed., Princeton, 1947.

4. L. S. Shapley, R. N. Snow, Solutions of the general two-person zero-sum game with a finite number of strategies (to be published).

5. J. Ville, Sur la théorie générale des jeux où intervient l'habilité des joueurs (in Traité du Calcul des Probabilités et de ses Applications by Émil Borel and collaborators, Vol. 2, No. 5), Paris, 1938.

6. A. Wald, Foundations of a general theory of sequential decision functions, Econometrica, Vol. 15 (1947), pp. 279-313.

LSS:rb